

A world without Cause

Benjamin James

December 24, 2025

Causality is usually treated as discovery; a hard-won insight wrested from a stubborn world. We are told that once upon a time humanity lived amid superstition and confusion, and then, through patience and rigor, uncovered the fact that things happen because other things happen first. This story flatters us, which is one reason it has survived so well. It presents causality as revelation rather than inheritance, as though it were pried loose from reality rather than quietly absorbed from the air we breathed while learning to speak. In truth, causality arrives far too early and far too smoothly to deserve that status. It is not discovered so much as assumed. It is not reasoned into existence but taken up as soon as explanation itself becomes possible.

Long before anyone is capable of asking whether causality is true, they are already using it. It enters alongside language, not as theory but as habit. This happened because of that. Don't do this or that will happen. We are not offered causality as a hypothesis to be tested, but as reassurance that the world is navigable and predictable. It promises that events are not arbitrary, that outcomes are tethered to actions, and that understanding consists in tracing a line between the two. This promise is generous and, at small scales, often accurate enough to be life-saving. If you touch fire, you will be burned. If you eat the wrong berry, you will be sick. These are not metaphysical commitments; they are survival heuristics. They work because the environments in which they evolved were narrow, slow, and forgiving of simplification.

This is how causality earned its authority. Not by correspondence to the deepest structure of reality, but by repeated success under constraint. It compresses experience into sequences that can be remembered, taught, and transmitted. It turns the overwhelming richness of our world into manageable stories with beginnings and ends. That compression is not a flaw; it is the entire point. Without it, culture would be impossible. Every generation would have to relearn everything from scratch, unable to pass on what worked because it

could not be summarized. Causality makes experience portable. It packages the past in a form that can be carried forward.

Causality resembles other great explanatory compressions we have since outgrown. Humoral theory organized medicine for centuries by reducing the body to four interacting fluids. Ether explained the propagation of light by filling the apparent emptiness of space with a medium that made intuitive sense. Divine will accounted for fortune and disaster alike by anchoring them to an intelligible agent. None of these were foolish ideas. They were powerful precisely because they simplified. They allowed people to act, coordinate, teach, and feel that the world was not mocking them with pure randomness. They endured until the environments they were meant to navigate outgrew them.

Causality has lasted longer than most because it is extraordinarily flexible. Unlike humors or ether, it does not specify much about what the world is made of. It is content with structure alone. There must be a before and an after. There must be a connection. There must be a reason. This minimalism is its strength. It can be applied almost anywhere, from physics to psychology to politics, without ever announcing its limits. It feels pre-theoretical because it is closer to grammar than to doctrine. We do not ask whether subject precedes verb before forming a sentence; we simply speak. In the same way, we do not ask whether causes precede effects before explaining an event; we simply explain.

Because causality functions grammatically, it becomes invisible. We notice it only when it breaks, and even then, we tend to treat the breakage as a local error rather than a structural mismatch. When predictions fail, we look for missing variables. When interventions misfire, we assume insufficient control. When explanations feel thin, we reach for more detail. Our underlying assumption, that the world itself is organized into causal chains waiting to be traced, remains untouched. It does not present itself as an assumption at all. It presents itself as what it would be irresponsible not to believe.

This is where our illusion of obviousness sets in. Causality feels self-evident not because it is true in some ultimate sense, but because so much of our social, institutional, and psychological infrastructure is built around it. Schools teach it. Laws depend on it. Scientific papers are structured to deliver it. To question causality begins to feel like questioning the

possibility of explanation itself, as though one were advocating silence or mysticism. That reaction is understandable, but it is also revealing. When an idea feels indistinguishable from our capacity to think at all, it is often because our thinking has been trained to move through it by default.

The success of causality has another consequence. It invites overconfidence. Because it works well in constrained settings, we are tempted to apply it everywhere. We treat it not as a local approximation but as universal law. This is the point at which a useful compression becomes distortion. The story that once allowed us to survive now begins to dictate how the world must behave in order to remain intelligible to us. When the world fails to comply, we experience not curiosity but irritation. Something must be wrong, we think, because our causal story is so good.

It is not the world without causality that should strike us as strange. It is the world with causality, the idea that reality obligingly arranges itself into clean sequences with identifiable origins and controllable outcomes. That is an extraordinary claim, and one that we rarely defend because we rarely notice we are making it. We take causality to be obvious precisely because it spares us the work of noticing how much must be ignored to keep it functioning. None of this requires us to denounce causality or pretend we can do without it entirely. A tool does not become useless simply because it is limited. The danger lies elsewhere. It lies in forgetting that causality is a story we tell about the world in order to move through it, not a transcript of how the world organizes itself in our absence. A story can be indispensable without being literal. In fact, our most indispensable stories often are not. They survive because they stabilize us, not because they describe everything that is happening.

To recognize causality as a successful mistake is not to sneer at it. It is to place it properly in our intellectual lineage. It is to see it as an adaptive inheritance, brilliant at what it was designed to do and increasingly strained by the environments to which we now apply it. Once this shift occurs, something subtle but important changes. We no longer feel compelled to defend causality as obviously real. We can acknowledge its utility without mistaking that utility for truth. That downgrade may feel minor, but it is decisive. It opens the possibility that

the world has been operating perfectly well all along, while our most trusted story about how it works has been quietly overstaying its welcome.

Once causality is no longer treated as sacred, it becomes possible to ask a surprisingly difficult question. What do we actually mean when we say that one thing causes another? The difficulty is not technical. It is that the answer feels too obvious to require articulation. Cause is assumed to be a basic feature of reality, like space or time, something so fundamental that attempting to define it would be redundant. Yet when pressed, what emerges is not a single clear commitment but a cluster of expectations that travel together under the same name. Causality, in practice, is not one idea. It is a toolkit.

At the most familiar level, causality means sequence. This happened, then that happened. This ordering itself carries explanatory weight. Earlier is treated as more important than later, antecedent as more real than consequent. Once this ordering is accepted, explanation becomes a matter of tracing a line backward until it reaches a satisfactory stopping point. This assumption is rarely examined because it aligns so neatly with how our stories are told and remembered, but sequence alone does not give us causality. It merely gives us chronology. The slide from one to the other happens so quickly that it often goes unnoticed. Layered onto sequence is the idea of intervention. If A caused B, then changing A should change B. This is where causality acquires its practical appeal. It promises leverage. The world is no longer just ordered; it is manipulable. Causes become handles, points where one can reach in and redirect the flow of events. This expectation is so central that it often defines what we mean by a “real” cause as opposed to a mere correlation. A cause is not just something that came before; it is something that could have been adjusted. This move quietly transforms explanation into control.

A third component completes the bundle, the counterfactual substitution. To say that A caused B is to imply that if A had not occurred, B would not have occurred either. This imagined alternative is treated as real, reachable, and informative. Counterfactuals allow us to reason about responsibility, regret, and policy. They give causality its moral and strategic depth. Without them, the past would feel fixed and opaque and the future unresponsive to

choice. With them, our world appears full of forks, each representing a decision that mattered.

Taken together, these three elements, sequence, intervention, and counterfactual substitution, form what most people mean by causality, whether or not they would articulate it this way. The crucial point is that this bundle is almost never defended explicitly. It is treated as background reality, our default framework within which thinking must occur. One can argue about which causes are correct, or which variables matter most, but rarely about whether this structure itself is optional. To question it feels like questioning the possibility of explanation rather than a particular explanatory style.

This is how causality masquerades as neutral description. It presents itself as simply reading off the structure of the world, when in fact it is importing a set of powerful commitments. The first is a metaphysics of control. To explain causally is to assume that the world is, at least in principle, steerable through localized adjustments. Even when we acknowledge complexity, we tend to treat it as a temporary obstacle to control rather than a feature that might resist our framing altogether. The language of causes encourages us to think in terms of levers and dials, even when the system in question behaves more like a shifting field than a machine. Alongside control comes moral legibility. Causality allows responsibility to be assigned cleanly. If A caused B, then A can be praised or blamed. This is enormously useful for legal systems, social coordination, and personal narratives of selfhood. It gives moral judgment a geometry. There is a line to follow and a point to indicate. Harm can be localized and credit can be claimed. Without causality, moral evaluation becomes diffuse and uncomfortable, spread across interactions and conditions rather than concentrated in a single act or actor. The attraction of causality here is obvious, but it is an attraction that should be named rather than smuggled in.

Predictive confidence completes our picture. If causes are stable and identifiable, then the future becomes, if not fully predictable, at least tractable. One can plan, optimize, and intervene. This confidence persists even when our predictions repeatedly fail, because failure is attributed to insufficient data or inadequate modeling rather than to a mismatch between our causal frame and the system being studied. Causality promises that with

enough refinement, the world will eventually yield. That promise is rarely scrutinized because it aligns so well with our institutional incentives and personal hopes.

None of this makes causality incoherent or useless. On the contrary, it explains why it is so difficult to dislodge. The problem is not that causality is wrong in some abstract sense, but that it does too much work while pretending to do very little. By presenting itself as a simple description of “how things happen,” it quietly installs a worldview in which control, blame, and prediction are treated as natural extensions of understanding. Once this installation is complete, alternatives begin to look like abdications rather than revisions.

Causality is best understood not as an ontological primitive, or a basic feature of reality, but as a narrative compression algorithm. It takes high-dimensional, interdependent processes and reduces them to sequences that can be told, taught, and acted upon. Like all compression schemes, it discards information. It smooths feedback, flattens simultaneity, and hides the degree to which outcomes depend on distributed conditions rather than isolated events. The resulting story is not false so much as selectively blind.

This explains both the power and the limits of causal reasoning. In environments where interactions are sparse, feedback is slow, and boundaries are clear, compression works remarkably well. Loss of information does not matter because the discarded details do not significantly affect outcomes. In such settings, causality feels not just useful but accurate. Our trouble begins when the same compression is applied to systems where interactions are dense, feedback is rapid, and boundaries are porous. There, discarded information is precisely what governs behavior.

Once causality is seen as a toolkit rather than a fact, it becomes possible to evaluate it without either reverence or contempt. One can ask where it works, where it fails, and what it hides. One can notice that the confidence it generates is often inversely proportional to the adaptability of the system being described. One can begin to separate explanation from control and narration from leverage without abandoning our ability to reason altogether. This shift does not require new metaphysics so much as a refusal to treat an old one as mandatory.

The downgrade is subtle but decisive. Causality no longer stands as an unquestioned backdrop against which all thinking occurs. It becomes one instrument among others, suitable for some tasks and poorly suited for others. We should not stop using causal language or deny its many successes. We should notice what is being assumed each time that language is used, and to recognize that those assumptions are optional. Causality remains available, but it is not inevitable. The world does not depend on it, even if our stories long have.

At this point our difficulty is no longer whether causality is useful, but when it is useful, and the answer is awkwardly specific. Causality excels at explaining what has already settled. It struggles, often invisibly, with what is still in motion. This is not a minor flaw or an edge case. It is structural. Causal reasoning operates most confidently after events have congealed into something stable enough to narrate. Before that moment, while systems are still adjusting, it offers far less guidance than its reputation suggests.

This claim can be stated plainly and without ceremony. *Causality explains the past far better than it governs the present or predicts the future.* This should not be controversial, but it rarely receives sustained attention. Instead, the success of causal explanation after the fact is quietly reinterpreted as evidence that the same framework must also have been operating prospectively. We mistake narrative closure for navigational power. The result is a persistent category error, repeated so often that it has become difficult to see.

Consider decision-making, not in theory but as it is lived. People regularly act before they can say why they acted. A choice presents itself as already made, and only then does the mind assemble a justification that makes the choice legible to itself and to others. These justifications are not fabrications. They point to real considerations, but they are arranged after the decision has already taken shape. Our causal story arrives once the system, cognitive, emotional, and contextual, has settled into an outcome. Our story feels explanatory precisely because it fits what is already there. It would have been much less helpful a moment earlier, when our decision was still forming.

This pattern repeats across domains with tedious consistency. Markets move before news. Prices shift, volatility spikes, capital flows rearrange themselves, and only afterward does

our explanatory machinery activate. Analysts identify signals, trace influences, and construct chains of causation that render the movement intelligible. These explanations are often sophisticated and internally consistent. They are also retrospective. They explain why the movement made sense, not how to anticipate it with the same clarity. When prediction fails, the failure is attributed to missing data or unforeseen shocks, rather than to the possibility that our causal frame itself lags behind the system it is describing.

Crises follow the same script. Before a collapse, uncertainty dominates. Signals are noisy, interpretations conflict, and warnings compete with reassurances. After collapse, clarity emerges almost instantly. The causes are obvious now, we say. The signs were all there. The story can be told cleanly, with a beginning, a middle, and an end. The same facts that produced ambiguity beforehand now produce inevitability. This transformation is not evidence that causality has finally revealed itself. It is evidence that once outcomes stabilize, they exert a powerful gravitational pull toward explanation. The past reorganizes itself around what narratively happened.

What deserves emphasis here is that causal explanations in these cases are rarely false. They are not hallucinations. They identify real relationships, real pressures, and real constraints. Their failure is temporal, not logical. They arrive after the window in which they could have guided action has closed. This lateness is easy to overlook because explanation feels like mastery. To understand why something happened creates a sensation of having understood how it works. That sensation is projected backward, as though understanding had been available all along.

Hindsight performs this projection automatically. Once an outcome is known, alternative possibilities collapse. Paths that were once live options are retroactively downgraded to noise. What could have happened becomes what never really could have happened. Causal narratives thrive in this environment because they benefit from a pruning that hindsight performs for free. Our story no longer has to account for uncertainty, hesitation, or partial information. It can proceed as though the correct path was always there, waiting to be discovered.

This is how contingency becomes inevitability. Our explanation does not lie; it simplifies. It removes the instability that defined the situation at the time by reconstructing a past from the vantage point of our present. The more detailed our explanation, the stronger this effect becomes. Each added detail makes the outcome feel more anchored, necessary, and caused. We forget that the explanation could not have been delivered with the same confidence before the outcome occurred.

The authority of causal narratives is therefore time-dependent. They gain force as systems settle. Once variables stop shifting and feedback loops slow down, causal attribution becomes easier and more persuasive. This is precisely when causal reasoning looks most impressive. It is also precisely when it is least useful for guiding what comes next. Our clarity is earned by the system's stabilization, not by the explanatory framework itself.

This distinction matters because it exposes a quiet slippage. We treat the success of causal explanation after the fact as evidence that the same reasoning can be used prospectively with sufficient refinement. We invest in better models, more data, and faster computation, assuming that the remaining gap is technical rather than conceptual. When predictions fail again, and again, we repeat the process. Our underlying assumption, that causality ought to work forward in time if only we could see clearly enough, remains untouched.

The pattern does not change. Explanation continues to outrun prediction. Control continues to lag behind understanding. The world continues to surprise us in ways that become immediately obvious once surprise is no longer possible. At some point repetition should prompt reconsideration. The failure is too consistent to be dismissed as noise.

The problem is not that causality cannot say anything about the future. It can, in constrained environments with limited feedback and slow dynamics. The problem is that we generalize from those successes to domains where the conditions that make causal reasoning prospective simply do not hold. In adaptive systems, the present is not a fixed platform from which the future can be projected. It is a moving boundary shaped by our very attempts to understand and influence it.

This is why causal narratives feel most confident at precisely the moment when their practical relevance diminishes. Once a system has settled, explanation flourishes. Before it

settles, while action still matters, explanation hesitates. The mismatch is not accidental. Causality requires stable relations to operate cleanly. Stability, by definition, arrives late. Impatience here is justified. This is not a subtle paradox or a technical corner case. It is a basic timing issue that has been obscured by habit. We have allowed the authority of hindsight to stand in for foresight. We have treated stories about our past as evidence of control over our future. In doing so, we have repeatedly confused narration with navigation. Once this confusion is named, it becomes difficult to unsee. Explanation no longer masquerades as prediction. Understanding no longer automatically implies leverage. I am not asking to abandon causal reasoning, only to place it where it belongs, as a powerful retrospective tool that organizes what has already happened, not as a reliable guide for what is still unfolding. That distinction, once drawn cleanly, removes much of the mystique surrounding causal certainty. What remains is not ignorance, but a clearer sense of what explanation can and cannot do in real-time.

Intervention begins with a picture that feels almost too reasonable to question. A system is imagined as a collection of parts, each with a role, and each separable in principle from the others. To intervene is to reach in, adjust one part, and withdraw, leaving the rest unchanged. This picture is not childish; it is inherited from domains where it works well enough to inspire confidence. Machines tolerate this treatment. Controlled experiments depend on it. Even social rituals sometimes reward the pretense. Our problem is not that this picture is naïve, but that it is taken to be general.

In lived systems, the moment of intervention is not a moment of contact but a moment of transformation. The system does not pause while an adjustment is made. It does not hold still so that the intervention can remain local. The act of intervening becomes part of what the system must now respond to. Expectations shift, incentives realign, and attention redistributes. What looks like a bounded change reveals itself as a signal that propagates outward, altering the conditions under which everything else operates. The imagined distinction between our intervention and the system dissolves immediately.

This is the sense in which intervention fails structurally rather than accidentally. It is not that the wrong lever was pulled or the wrong variable adjusted. It is the assumption a stable lever

was misplaced. The system one intended to modify no longer exists once modification is attempted. This central line bears repeating because it is so often ignored. *The moment you intervene; the system you intervened in no longer exists.* This is not poetic license. It is a practical description of what happens when adaptive systems are treated as if they are static.

Resistance is a misnomer. When policies produce unintended consequences, when treatments generate side effects, or when reforms trigger evasive behavior, our language of resistance suggests obstinacy or failure. The system is said to be pushing back, but what is being labeled resistance is simply adaptation. The system is doing what it does best, reconfiguring itself in response to new constraints. Disappointment arises not because the system behaved strangely, but because our intervention was designed as though adaptation were optional.

This pattern is visible wherever interventions are attempted at scale. Economic incentives reshape behavior before they achieve their stated goals. Regulations alter the landscape they are meant to stabilize by changing what is profitable, visible, or punishable. Educational reforms modify how success is signaled long before they modify how learning occurs. Each time, the response is framed as complexity getting in the way of a simple plan. Rarely is the plan itself questioned at the level of its underlying assumptions.

Counterfactual reasoning enters at this point, offering comfort where intervention falters. If the intervention did not work, it must be because it was poorly chosen, poorly timed, or insufficiently aggressive. One imagines an alternative world in which the right change was made at the right moment, producing the desired outcome. This imagined alternative serves several functions at once. It preserves our basic model; it allows responsibility to be assigned or absolved; and it restores our sense that the system was, in principle, controllable.

Counterfactuals are rarely malicious. They are not invented to deceive. They arise naturally as a way of managing regret and preserving agency. To say “if only we had done X” is to assert that the world remains navigable and that outcomes are still tethered to choices in a way

that makes moral sense. Without this tether, failure feels arbitrary and success feels unearned. Counterfactuals reassure us that the story still has joints.

In policy contexts, this reassurance becomes institutionalized. After-action reports trace alternate paths with impressive clarity. They identify missed opportunities, unpulled levers, and decisions that could have been different. These analyses are often careful, data-rich, and internally coherent. They are also deeply dependent on the very outcome they are explaining. The plausibility of the counterfactual world is borrowed from the actual one. The conditions that make the alternative seem reachable are precisely those that were produced by the path that was taken.

This dependence is easy to miss because counterfactuals are constructed retrospectively. They are shaped by knowledge that was unavailable at the time of action. Signals that were ambiguous become decisive, and constraints that were invisible become obvious. Our imagined alternative is smoothed of uncertainty and friction, presented as a simple substitution rather than a reconfiguration of the entire field. What is rarely acknowledged is that many of these alternatives were not merely unlikely; they were structurally unavailable. The system could not have occupied those states without first becoming something else.

At a personal level, the same logic applies. Regret relies on the belief that a different choice would have led to a different outcome, all else remaining equal, but all else never remains equal. The confidence required to make the imagined choice often depends on experiences that only occurred because the original choice was made. The insight one wishes one had possessed earlier was produced by the consequences of not possessing it. Our alternative-self imagined in regret is a hybrid, borrowing capacities from a future that only exists because our past unfolded as it did.

Intervention and counterfactuals thus collapse into a single failure mode. Both assume that systems can be locally modified without globally reconfiguring, and that alternate outcomes can be reached by adjusting isolated inputs. Both depend on a picture of possibility as a branching tree, where each decision point offers a menu of equivalent futures differing only in the choice made. In adaptive systems, this picture does not hold. Possibility is not evenly

distributed across branches; it is shaped by accumulated constraints, histories, and feedback.

A more accurate image is not a tree but a terrain. Movement is guided by slopes, ridges, and basins. Some directions are easy, others steep, others effectively inaccessible. Paths that look adjacent on a map may be separated by barriers that only become visible in motion. Small changes can redirect trajectories if they align with existing gradients, and massive efforts can fail if they push against them. Importantly, the terrain itself changes as movement occurs. Each step reshapes the ground that the next step must traverse.

Seen from this perspective, the failure of intervention is no longer mysterious, and the allure of counterfactuals loses much of its force. There is no single point at which the right lever could have been pulled, because there was no stable lever to begin with. There are no cleanly reachable alternate worlds waiting just offstage, because the worlds that feel reachable are defined by the paths already taken. Possibility is not a set of branches but a topology that evolves under pressure.

This does not mean action is futile or responsibility evaporates. It means that action must be understood as participation rather than control, and responsibility is distributed rather than localized. The damage caused by intervention and counterfactual reasoning lies not in their use, but in their overextension. When they are treated as universal rather than situational, they generate confidence where humility would be more appropriate and blame where understanding is required.

By the end of this line of reasoning, lever-based thinking begins to feel brittle, and branch-based imagination begins to feel artificial. Our world does not present itself as a machine waiting to be tuned or a decision tree waiting to be explored. It presents itself as a shifting landscape in which outcomes emerge from alignment and persistence rather than from isolated acts. To continue insisting on clean interventions and reachable counterfactuals in such a landscape is not to be rigorous, but to be out of step with how living systems actually behave.

What remains when causality is set aside is not disorder, but a different way of describing order. One that was always available but rarely taken seriously because it does not flatter our

habits of explanation. The world does not dissolve into noise when we stop asking what caused what. It continues to hold together, change, surprise, stabilize, fail, and recover. What disappears is not intelligibility, but a particular style of intelligibility that insists on origins, levers, and clean sequences. In its place appears something less dramatic and more faithful. A picture of reality as navigation through constraint, as the temporary holding of form under pressure, as continuous reconfiguration without authorship.

Stability is no longer something produced by a cause; it is something maintained. Patterns persist not because they were initiated correctly, but because the conditions that sustain them continue to align. When those conditions drift, the pattern drifts with them. There is no single moment where one can point and say, “this is where it began,” because beginnings are themselves emergent features of later narration. Systems settle into forms that can hold, and those forms last exactly as long as the pressures acting on them allow. Nothing mystical is required to see this. One simply stops insisting that persistence must be explained by reference to a privileged event in the past.

Constraint does most of the work that causality once claimed. Not constraint as prohibition, but constraint as shape. What can happen is bounded by what is already in place. By material limits, institutional arrangements, habits, expectations, histories, and feedback that do not announce themselves as causes but quietly channel movement. Within these bounds, adaptation occurs continuously. Systems adjust not toward a goal, but toward viability. They do not seek optimal outcomes; they seek states that can be sustained under current conditions. When those conditions shift, so does the definition of what counts as viable.

This is why change so often looks incremental until it suddenly does not. The underlying structure has been adjusting all along, accumulating tension, and redistributing load, until a configuration that once held no longer can. The transition is then described as a rupture, turning point, or a cause finally making itself known. In reality, the system simply moved from one temporary stability to another. The drama belongs to our story, not the process.

Described this way, reality does not need single origins. It does not require privileged levers. It does not assign authorship to outcomes in the way our narratives prefer. Things happen

because they can, because they fit, and because they are sustained long enough to be noticed. When they stop fitting, they dissolve or transform. This does not make outcomes arbitrary. It makes them conditional in a richer sense than causality allows. They depend not on one thing, but on many things continuing to cohere.

Ethics, once freed from the arrow, becomes more exacting rather than more permissive. Responsibility no longer collapses onto isolated decisions or individuals, but distributes itself across structures, roles, and sustained patterns of participation. This is uncomfortable for those who prefer moral clarity delivered in the form of verdicts. It is also more accurate. Harm rarely persists because of a single act. It persists because conditions are maintained that allow harm to continue. Praise, too, becomes less intoxicating. Success is harder to attribute to personal virtue when it is clearly scaffolded by circumstances, support, timing, and alignment that no one fully controls.

Punishment and praise do not disappear in this view, but they change function. They lose their theatrical finality and gain a more modest role as signals within a system. Instead of closing moral accounts, they become part of the ongoing process of adjustment. The question shifts from “who caused this?” to “what is being reinforced?” and “what needs to change for this pattern not to repeat?” Ethics moves from judgment to maintenance and from spectacle to repair.

This shift does not absolve individuals of agency. It reframes agency as situated rather than sovereign. People still act and choices still matter, but they matter as contributions to patterns rather than as origin points. Agency becomes the capacity to respond within constraint, to notice when a configuration is becoming brittle, and to participate in re-alignment rather than to imagine oneself standing outside the system issuing commands. This is a thinner notion of control and a thicker notion of responsibility.

Fear, at this juncture, is predictable. Without causality, do we lose our ability to teach? To govern? To hold anyone accountable? These fears explain much of causality’s persistence. It is pedagogically efficient. It makes institutions legible. It comforts individuals by assuring them that effort maps cleanly onto outcome. To loosen causality feels, at first, like loosening agency itself.

This is another inherited confusion. Agency does not require causality to be literal; it requires that actions make a difference. In adaptive systems, actions make a difference all the time; they just do so by altering conditions rather than producing outcomes directly. Teaching still works, but it is understood as shaping environments of learning rather than transmitting causes of understanding. Governance still functions, but it is recognized as the art of maintaining viable constraints rather than pulling policy levers. Accountability still matters, but it is oriented toward sustained participation in harmful or beneficial structures rather than isolated blame.

Why, then, do we cling so tightly to causal language? Partly because it is easy to teach. Partly because it fits neatly into forms, laws, and models that demand clarity at the expense of fidelity. Partly because it soothes our anxiety by making the world feel controllable. Partly because abandoning it sounds like surrender. We have learned to equate the arrow with agency, the cause with power. To give that up feels like retreat.

It is not retreat. It is repositioning. We need not ban causal talk or pretend it has no place. It is to stop treating it as literal. To recognize it as one story among others; an old, successful story that helped us survive in narrower worlds and now misleads us in wider ones. Causality can still be used as a heuristic, convenience, or shorthand without being mistaken for the structure of reality itself. Nothing dramatic follows from this recognition. There is no revolution in thought, no sudden mastery of complexity. There is only a subtle shift in what we expect explanations to do.

We stop demanding they deliver control. We stop mistaking clarity for power. We allow ourselves to describe systems as they are rather than as we wish them to be. We loosen our grip on causality without throwing it away. We treat it as a story that once helped us coordinate and now obscures more than it reveals when taken too seriously. We become comfortable describing persistence without authorship, change without beginnings, and responsibility without villains. The world does not resist this description. It has been operating this way all along.

Reality never promised to behave causally. We made that promise to ourselves, and for a long time it kept us oriented. It may still have its uses, but we are no longer required to believe

that its promise was ever written into the fabric of things. We simply learned to speak as if it were.

References

James, B. (2025). From Causation to Coherence: Deconstructing Pearl and Reconstructing Reason. PhilPapers. <https://philpapers.org/rec/JAMFCT>

James, B. (2025). The Cult of Causality, Pearl's priesthood has no clothes. PhilPapers. <https://philpapers.org/rec/JAMTCO-35>